

WHAT IS CLAIMED IS:

1. A method of scheduling high-priority packets in a metro Ethernet switch, the method comprising the steps of:

determining a maximum queuing delay allowed for at least two high-priority packets in a queue in the switch;

determining which one of the at least two high-priority packets has the smallest maximum queuing delay allowed; and

scheduling the one of the at least two high-priority packets determined to have the smallest maximum queuing delay allowed before the remaining ones of the at least two high-priority packets.

2. The method of claim 1 wherein the step of determining the maximum queuing delay allowed comprises the steps of:

creating a POS table that lists, for each high-priority packet that has entered the switch, a position of the high-priority packet in a queue of the switch, a time the high-priority packet entered the queue, and an intended destination of the high-priority packet;

creating a Qmax table for storing a maximum allowed queuing delay for each of several possible intended destinations; and

using the Qmax table and the POS table to determine the maximum queuing delay allowed for each of the high-priority packets in the queue of the switch.

3. The method of claim 2 wherein the step of creating a Qmax table comprises the steps of, for each label switched path ("LSP") between the switch and one of the possible intended destinations:

determining a number of hops along the LSP; and

dividing a maximum queuing delay allowed for the LSP by the number of hops along the LSP to determine the maximum queuing delay allowed for each hop.

4. The method of claim 2 wherein the step of creating a Qmax table is performed only once during LSP setup.

5. The method of claim 2 further comprising the step of updating the POS table each time a new high-priority packet enters the queue.

6. The method of claim 1 wherein the steps of determining a maximum queuing delay allowed, determining which one of the at least two high-priority packets has the smallest maximum queuing delay allowed, and scheduling the one of the at least two high-priority packets determined to have the smallest maximum queuing delay allowed before the remaining ones of the at least two high-priority packets are performed each time a new high-priority packet enters the queue.

7. The method of claim 1 wherein the queue is capable of performing an n -packet look-ahead.

8. A method of scheduling high-priority packets in a metro Ethernet switch, the method comprising the steps of:

creating a first table that lists, for each high-priority packet that has entered the switch, a position of the high-priority packet in a queue of the switch, a time the high-priority packet entered the queue, and an intended destination of the high-priority packet;

creating a second table for storing a maximum allowed queuing delay for each of several possible intended destinations; and

using the first and second tables to determine the maximum queuing delay allowed for each of the high-priority packets in the queue of the switch.

9. The method of claim 8 further comprising the step of:

determining a maximum queuing delay allowed for at least two high-priority packets in a queue in the switch.

10. The method of claim 9 further comprising the step of:

determining which one of the at least two high-priority packets has the smallest maximum queuing delay allowed.

11. The method of claim 10 further comprising the step of:

scheduling the one of the at least two high-priority packets determined to have the smallest maximum queuing delay allowed before the remaining ones of the at least two high-priority packets.

12. The method of claim 9 wherein the step of determining a maximum queuing delay allowed for at least two high-priority packets in a queue in the switch is performed with reference to the first and second tables.

13. The method of claim 8 wherein the step of creating a second table comprises the steps of, for each label switched path ("LSP") between the switch and one of the possible intended destinations:

determining a number of hops along the LSP; and

dividing a maximum queuing delay allowed for the LSP by the number of hops along the LSP to determine the maximum queuing delay allowed for each hop.

14. The method of claim 8 wherein the step of creating a second table is performed only once during LSP setup.

15. The method of claim 8 further comprising the step of updating the first table each time a new high-priority packet enters the queue.

16. The method of claim 8 wherein the queue is capable of performing an n -packet look-ahead.

17. Apparatus for scheduling high-priority packets in a metro Ethernet switch, the apparatus comprising:

means for determining a maximum queuing delay allowed for at least two high-priority packets in a queue in the switch;

means for determining which one of the at least two high-priority packets has the smallest maximum queuing delay allowed; and

means for scheduling the one of the at least two high-priority packets determined to have the smallest maximum queuing delay allowed before the remaining ones of the at least two high-priority packets.

18. The apparatus of claim 17 wherein the means for determining the maximum queuing delay allowed comprises:

a POS table that lists, for each high-priority packet that has entered the switch, a position of the high-priority packet in a queue of the switch, a time the high-priority packet entered the queue, and an intended destination of the high-priority packet;

means for creating a Qmax table for storing a maximum allowed queuing delay for each of several possible intended destinations; and

means for using the Qmax table and the POS table to determine the maximum queuing delay allowed for each of the high-priority packets in the queue of the switch.

19. The apparatus of claim 18 wherein the means for creating a Qmax table comprises, for each label switched path ("LSP"):

means for determining a number of hops along the LSP; and

means for dividing a maximum queuing delay allowed for the LSP by the number of hops along the LSP to determine the maximum queuing delay allowed for each hop.

20. The apparatus of claim 18 wherein the Qmax table is created during LSP setup.

21. The apparatus of claim 18 further comprising means for updating the POS table each time a new high-priority packet enters the queue.

22. The apparatus of claim 17 wherein the queue is capable of performing an n -packet look-ahead.

23. A packet switch comprising:

a queue containing a plurality of packets received at the switch; and

a scheduler for scheduling transmission of the packets in the queue, wherein when the queue contains at least two high-priority packets, the scheduler schedules the one of the at least two high-priority packets determined to have a smallest maximum queuing delay allowed before the remaining ones of the at least two high-priority packets.

24. The packet switch of claim 23 further comprising a state machine for:

maintaining a POS table that lists, for each high-priority packet that has entered the switch, a position of the high-priority packet in a queue of the switch, a time the high-priority packet entered the queue, and an intended destination of the high-priority packet; and

maintaining a Qmax table for storing a maximum allowed queuing delay for each of several possible intended destinations.

25. The packet switch of claim 24 wherein the scheduler uses the Qmax table and the POS table to determine the maximum queuing delay allowed for each of the high-priority packets in the queue of the switch.

26. The packet switch of claim 24 wherein the maximum allowed queuing delay for each of several possible intended destinations is determined by determining a number of hops along the LSP and dividing a maximum queuing delay allowed for the LSP by the number of hops along the LSP to determine the maximum queuing delay allowed for each hop.